

Signal Detection Theory

Professor David Heeger

November 12, 1997

The starting point for signal detection theory is that nearly all decision making takes place in the presence of some uncertainty. Signal detection theory provides a precise language and graphic notation for analyzing decision making in the presence of uncertainty.

Simple Forced Choice

I begin here with the classic example of detecting brief, dim flashes of light in a dark room. Imagine that we use a simple forced-choice method in which the light is flashed on half of the trials (randomly interleaved). On each trial, the subject must respond "yes" or "no" to indicate whether or not they think the light was flashed. We assume that the subjects' performance is determined by the number of photon absorptions/photopigment isomerizations on each trial.

There are two kinds of noise factors that limit the subject's performance: *internal noise* and *external noise*.

External noise. There are many possible sources of external noise. The main source of external noise to consider here is the quantal nature of light. On average, the light source is set up to deliver a certain stimulus intensity, say 100 photons. A given trial, however, there will rarely be exactly 100 photons emitted. Instead, the photon count will vary from trial to trial following a Poisson distribution.

Internal noise. Internal noise refers to the fact that neural responses would be noisy, even if the stimulus was exactly the same on each trial. Some of the emitted photons will be scattered by the cornea, the lens, and the other goopy stuff in the eye. The number of scattered photons will vary randomly from trial to trial. Of the photons that reach the photoreceptors, not all of them will be absorbed by the photopigments. There are other sources of internal (neural) noise as well, but I will ignore those for the time being.

In practice, of course, it would be impossible to measure the number of photons absorbed on any given trial because we would have to record simultaneously from all the rods in the retina. It is, however, possible to characterize the probability that a certain number of photons will be absorbed. Each of the relevant factors (number of photons emitted, number of photons scattered, number of photons absorbed) can be modeled as a Poisson process. A sequence of Poisson processes behaves altogether like a single Poisson process with an overall rate constant equal to the products of all of individual rate constants. For example, assume that 100 photons are emitted on average, that 10% of those photons pass through the eyes' optics on average, and that 10% of those are absorbed by photoreceptors on average. Then there will then be 1 photon absorbed on average for each trial on which the light is flashed, and this number will vary from trial to trial following a Poisson distribution.

On trials for which no light is flashed, there will still typically be some non-zero level of response, due to thermal isomerizations of photopigment molecules. Barlow called this the "dark light" because a spontaneous isomerization will lead to the same neural signal as if a photon was actually absorbed. The subject will not be able to tell the difference between real light and dark light.

Internal response probability density functions. Because the task is so hard, there is always some uncertainty as to what was there or not. Either there was a flash (signal plus noise) or there was no flash (noise alone). Either the subject saw the flash (they respond "yes") or they did not (they respond "no"). There are four possible outcomes: hit (signal present and subject says "yes"), miss (signal present and subject says "no"), false alarm (signal absent and subject says "yes"), and correct rejection (signal absent and subject says "no"). Hits and correct rejections are good. False alarms and misses are bad.

Figure 1 shows a graph of two hypothetical internal response curves. The curve on the left is for the noise-alone trials, and the curve on the right is for the signal-plus-noise trials. The height of each curve represents how often that level of internal response will occur. On noise-alone trials, in this example, there will generally be about 10 units of internal response (i.e., 10 photopigment isomerizations). However, there will be some trials with more (or less) internal response because of the internal and external noise.

Notice that we never lose the noise. The internal response for the signal-plus-noise case is generally greater but there is still a distribution (a spread) of possible responses. Notice also that the curves overlap, that is, the internal response for a noise-alone trial may exceed the internal response for a signal-plus-noise trial.

The role of the criterion. There are two main components to the decision-making process: *stimulus strength* and criterion. The stimulus strength affects the probability density functions in the obvious way: a stronger signal (brighter flash) will shift the signal-plus-noise curve to the right. More on this later.

The second component of the decision process is quite different. The subject is being

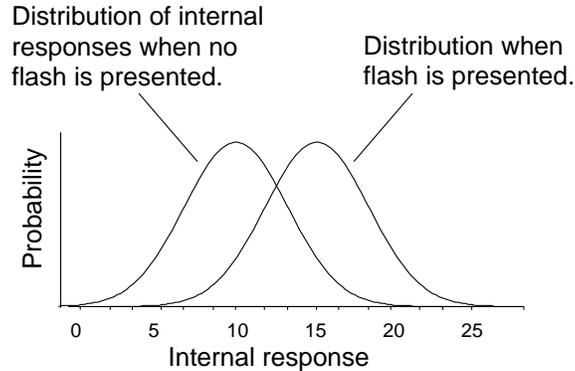


Figure 1: Internal response probability density functions for noise-alone and for signal-plus-noise trials.

asked to use their own judgement in making a decision. Different subjects may feel that the different types of errors are not equal. You might even emphasize to one subject that they should be careful not to make any false alarms, or you might offer to pay them more for getting more hits.

Perhaps the simplest strategy that the subject can adopt is to pick a *criterion* location along the internal response axis. Whenever the internal response is greater than this criterion they respond “yes”. Whenever the internal response is less than this criterion they respond “no”.

An example criterion is indicated by the vertical lines in Figure 2. The criterion line divides the graph into four sections that correspond to: hits, misses, false alarms, and correct rejections. On both hits and false alarms, the internal response is greater than the criterion, because the subject is responding “yes”. Hits correspond to signal-plus-noise trials when the internal response is greater than criterion, as indicated in the figure. False alarms correspond to noise-alone trials when the internal response is greater than criterion, as indicated in the figure.

Suppose that the subject chooses a low criterion (Figure 3, top), so that they respond “yes” to almost everything. Then they will never miss a flash when it is present and they will therefore have a very high hit rate. On the other hand, saying “yes” to almost everything will greatly increase the number of false alarms. Thus, there is a clear cost to increasing the number of hits, and that cost is paid in terms of false alarms. If the subject chooses a high criterion (Figure 3, bottom) then they respond “no” to almost everything. They will rarely make a false alarm, but they will also miss many real flashes.

There is no way that the subject can set their criterion to achieve only hits and no false alarms; it is *inevitable* that some mistakes will be made. Because of the noise it is simply a true, undeniable, fact that the internal responses on noise-alone trials may exceed the internal responses on signal-plus-noise trials, in some instances. Thus the subject cannot always be right. They can adjust the *kind* of errors that they make by manipulating their criterion, the one part of this diagram that is under their control.

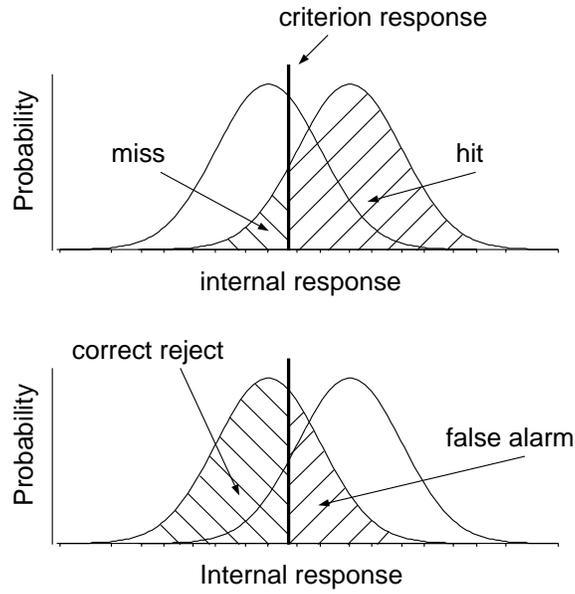


Figure 2: Internal response probability density functions for noise-alone and signal-plus-noise trials. Since the curves overlap, the internal response for a noise-alone trial may exceed the internal response for a signal-plus-noise trial. Vertical lines correspond to the criterion response.

The receiver operating characteristic. We can describe the full range of the subject's options in a single curve, called an *ROC curve*, which stands for receiver-operating characteristic. The receiver-operating characteristic captures, in a single graph, the various alternatives that are available to the subject as they move their criterion to higher and lower levels.

ROC curves (Figure 4) are plotted with the false alarm rate on the horizontal axis and the hit rate on the vertical axis. We already know that if the criterion is high, then both the false alarm rate and the hit rate will be very low. If we move the criterion lower, then the hit rate and the false alarm rate both increase. So the full ROC curve has an upward sloping shape. Notice also that for any reasonable choice of criterion, the hit rate is always larger than the false alarm rate, so the ROC curve is bowed upward. The subject may set the criterion anywhere, but any choice that they make will land them with a hit and false alarm rate somewhere on the ROC curve.

The role of signal strength. If we present a brighter flash (e.g., with 200 photons emitted per flash on average rather than 100), then the subject's internal response strength will, on the average, be stronger. Pictorially, this will have the effect of shifting the probability density function for signal-plus-noise trials to the right, a bit further away from the noise-alone probability density.

Figure 4 shows two sets of probability densities and two ROC curves. When the signal is stronger there is less overlap between the two probability density curves. When this

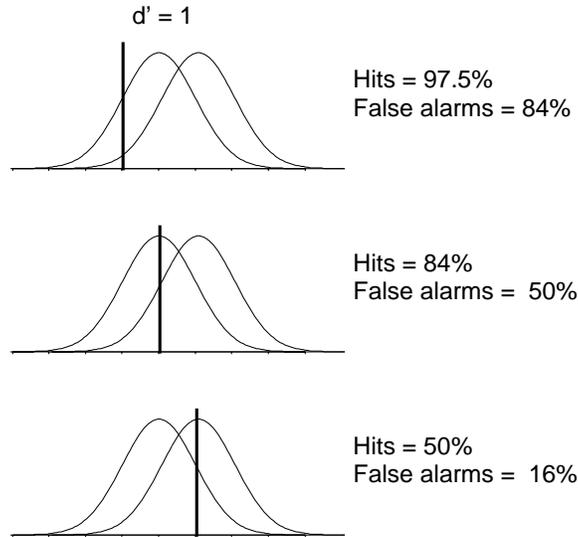


Figure 3: Effect of shifting the criterion.

happens the subject's choices are not so difficult as before. They can pick a criterion to get nearly a perfect hit rate with almost no false alarms. ROC curves for stronger signals bow out further than ROC curves for weaker signals.

Varying the noise. There is another aspect of the probability densities that also determines detectability: the *spread* of the curves. For example, consider the two sets of probability densities in Figure 5. The separation between the peaks is the same but the second set of curves are much skinnier. Clearly, the signal is much more discriminable when there is less spread (less noise) in the probability densities. So the subject would have an easier time setting their criterion in order to be right nearly all the time.

In our example, we have assumed Poisson noise so the absorption count variance is proportional to the mean absorption count. However, one can easily imagine situations in which the response variance depends on factors that are independent of the mean response.

Discriminability index (d'). Thus, the discriminability of a signal depends both on the separation and the spread of the noise-alone and signal-plus-noise curves. To write down a full description of how discriminable the signal is from no-signal, we want a formula that captures both the separation and the spread. The most widely used measure is called *d-prime* (d'), and its formula is simply:

$$d' = \frac{\text{separation}}{\text{spread}},$$

where the separation corresponds to the difference between the means, and the spread corresponds to the standard deviation of the probability densities. This number, d' , is

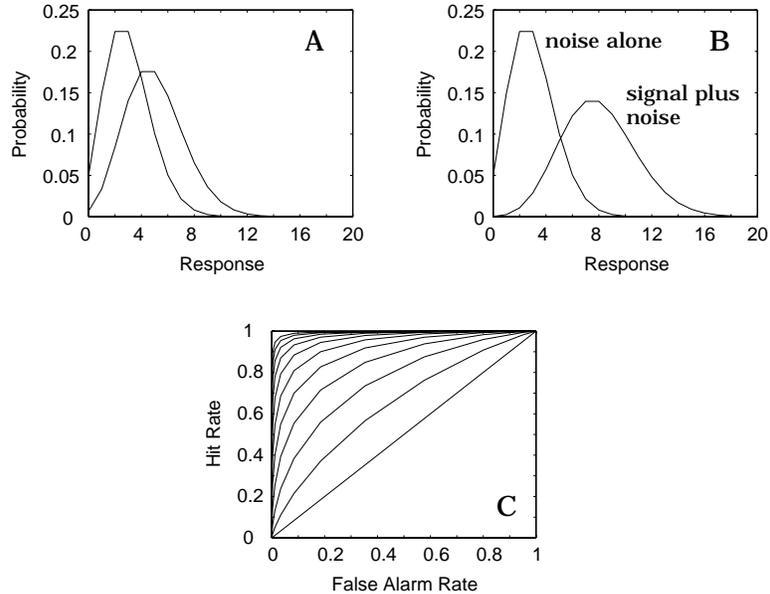


Figure 4: Internal response probability density functions and ROC curves for different signal strengths. When the signal is stronger there is less overlap in the probability of occurrence curves, and the ROC curve becomes more bowed. A: Probability density functions when the signal evokes an average of 2 photon absorptions per trial. B: Probability density functions when the signal evokes an average of 5 photon absorptions per trial. C: ROC curves for a series of signal strengths that evoke an average of $n = 0, 1, 2, \dots, 10$ photon absorptions per trial. In all cases the dark noise (average number of spontaneous isomerizations per trial) was 3.

a complete characterization of the detectability of the signal assuming that the noise follows a normal (Gaussian) distribution with a fixed variance, independent of the signal strength. This assumption of IID (independent and identically distributed) Gaussian noise is often reasonable approximations. However, if you have more information about the noise distribution (e.g., that it follows the Poisson distribution), you might as well use that information rather than assuming IID Gaussian noise.

The primary virtue of d' , and the reason that it is so widely used, is that its value does not depend upon the criterion the subject is adopting, but instead it is a true measure of the internal response.

Comparing neural responses with behavioral performance. Let's say that we carefully measure, in a separate experiment, the average number of spontaneous (thermal) isomerizations per trial. Then we can compute a series of ROC curves each corresponding to a different number of photon absorptions. Figure 4 shows such a family of ROC curves. Exactly how to compute these curves is illustrated in assignment3Tutorial.m.

Now we do our detection experiment in which we ask our subject to run 1000 trials. On half the trials, the flash is absent (noise-only trials) and on half the trials the light is

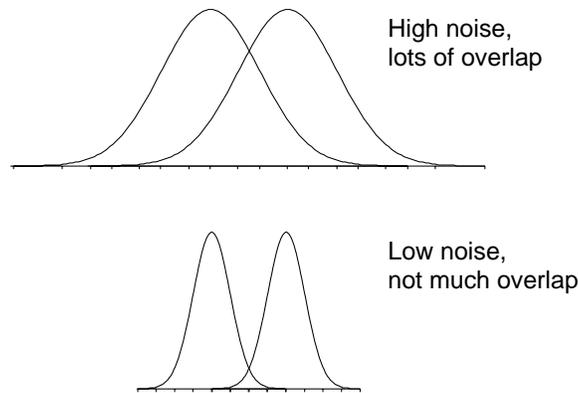


Figure 5: Internal response probability density functions for two different noise levels. When the noise is greater, the curves are wider (more spread) and there is more overlap.

flashed with a “fixed” intensity (e.g., with 100 photons emitted per flash on average). We count up the number of hits and false alarms, and that drops us somewhere on one of the ROC curves in Figure 4. Let’s say it lands us on the ROC curve labelled $n=1$ (1 photon absorption per trial, on average). Then we might conclude that people are capable of detecting a single photon of light. To check the assumptions of this analysis, we could double the intensity of the light so that 200 photons would be emitted per flash on average. Running the forced-choice experiment and counting the hits and false alarms in this case should drop us somewhere on the ROC curve labelled $n=2$. And so on for 300, 400, etc. emitted photons.

Notice that we need to know both the hit rate and the false alarm rate to get a measure of performance that is independent of the subject’s criterion.

Although nobody has actually done this experiment in exactly this way, there is quite a lot of evidence that people are capable of detecting a very dim flash of light for which only a few (maybe only 1, but probably more are necessary) photons are absorbed (Hecht *et al.*, 1942; Barlow and Levick, 1969; Barlow *et al.*, 1971; Sackitt, 1972). That is really quite surprising because it implies that a handful of photon absorptions give rise to a reliable neural signal that bubbles all the way up through the visual pathways. How do you think that might be possible, given the noisiness of cortical responses?

Two-Alternative Forced Choice

Another way to measure performance independent of any criterion is to use a two-alternative, forced-choice method. We’ll use the Newsome *et al.* direction discrimination experiments for an example. Newsome and colleagues (Newsome *et al.*, 1989; Salzman *et al.*, 1990, 1992; Britten *et al.*, 1992) recorded neural activity of MT neurons in response to stimuli consisting of a field of coherently moving dots superimposed on a field of randomly moving dots. The strength of the motion signal was controlled by varying the ratio of coherent to random dots. The coherent dots moved either in the cell’s preferred direc-

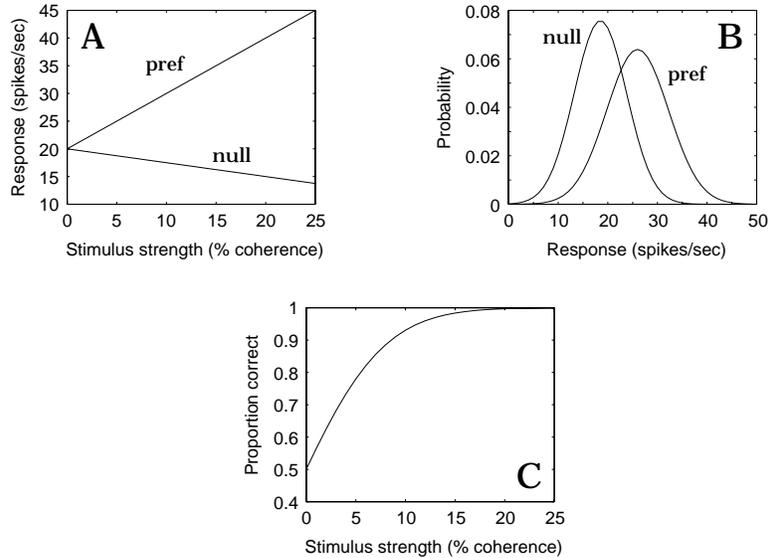


Figure 6: A: Simulated MT cell responses for preferred and null directions, as a function of motion coherence. B: Simulated response probability density functions for preferred and null directions, at 6% coherence. C: Simulated neurometric function (percent correct versus motion coherence).

tion or in the opposite (null) direction. The monkeys were trained to report the direction of motion by making an eye movement at the end of each trial.

Although Newsome *et al.* did not do it exactly this way, it will simplify matters if we have two intervals in each trial. A stimulus moves to the right on one interval (chosen randomly), and it moves to the left on the other interval. On each trial, the monkey must choose the interval during which the motion was rightward. Because the stimulus is optimized for the recorded neuron (covers the receptive field, moves in the preferred direction, etc.), one might hypothesize that the monkey monitors the response of that one neuron to make his decision, and chooses the interval that evokes the greater response. The task is very difficult for low coherence levels (e.g., below 5%), and very easy for high coherence levels (e.g., above 20%).

Figure 6 illustrates, via a simulation, how behavioral performance could depend on the noisy responses of a single MT neuron. Figure 6A plots a simulated example of how the mean response varies with coherence. For motion in the preferred direction, the simulated mean firing rate rises linearly with coherence. For motion in the null direction, the mean firing rate declines linearly with coherence. Figure 6B plot probability density functions for the responses to preferred and null directions at 6% coherence. In these simulations, the PDFs are normally distributed. The response variance was set equal to 1.5 times the mean response, consistent with a number of studies of the variability of cortical responses (e.g., Dean, 1981; Tolhurst *et al.*, 1983; Bradley *et al.*, 1987; Snowden *et al.*, 1992; Britten *et al.*, 1993; Softky and Koch, 1993). Figure 6C plots a neurometric function (percent correct versus motion coherence) computed from the responses of this simulated MT neuron.

| | |
|------------|---|
| c | Stimulus motion coherence |
| $r_p(c)$ | Response to preferred direction motion with coherence c . |
| $r_n(c)$ | Response to null direction motion. |
| $f_p(r c)$ | Response probability density function for preferred direction motion with coherence c . |
| $f_n(r c)$ | Response probability density function for null direction motion. |
| $F_n(r c)$ | Cumulative distribution for null direction motion. |

Table 1: Table of Mathematical Symbols.

To understand this last step of the simulation, we need to do a simple derivation using the mathematical notation introduced in Table 1. The probability of being correct on a given trial is given by:

$$P(\text{correct}) = P[r_p(c) > r_n(c)],$$

where $r_p(c)$ and $r_n(c)$ are random variables that are drawn from the response probability density functions, $f(r_p)$ and $f(r_n)$, for preferred and null motion directions, respectively. In words, the probability of being correct is given by the probability that a preferred direction stimulus will evoke a bigger response than a null direction stimulus. This probability depends on the two response probability densities:

$$P[r_p(c) > r_n(c)] = \int_0^\infty f_p(r|c) \left[\int_0^r f_n(r'|c) dr' \right] dr.$$

The integral in brackets represents the probability that the null direction response will be smaller than some criterion response r . This is multiplied by the probability that the preferred direction stimulus will evoke that criterion response, and the whole thing is summed/integrated over all possible criterion response levels. This equation can be simplified by noting that the integral in brackets is the cumulative distribution function for null direction motions. Hence,

$$P(\text{correct}) = \int_0^\infty f_p(r|c) F_n(r|c) dr.$$

The Newsome *et al.* experiments were actually done with one interval on each trial, in which the monkey made a forced choice judgement about the direction of motion. To analyze the data and compute the neurometric function, they assumed that the monkey monitored the responses of two neurons, the one that they were recording from and a second (that they called the “anti-neuron”) that was identical in all respects except that it had the opposite direction preference. They assumed that the monkey picked the direction on each trial corresponding to the neuron/anti-neuron that gave the larger response. This is formally identical to the two-interval forced choice case treated above.

References

- [1] H B Barlow and W R Levick. Three factors limiting the reliable detection of light by retinal ganglion cells of the cat. *Journal of Physiology (London)*, 200:1–24, 1969.

- [2] H B Barlow, W R Levick, and M Yoon. Responses to single quanta of light in retinal ganglion cells of the cat. *Vis. Res. Suppl.*, 3:87–101, 1971.
- [3] A Bradley, B Skottun, I Ohzawa, G Sclar, and R D Freeman. Visual orientation and spatial frequency discrimination: A comparison of single neurons and behavior. *Journal of Neurophysiology*, 57:755–772, 1987.
- [4] K H Britten, M N Shadlen, W T Newsome, and J A Movshon. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12:4745–4765, 1992.
- [5] K H Britten, M N Shadlen, W T Newsome, and J A Movshon. Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience*, 10:1157–1169, 1993.
- [6] A F Dean. The variability of discharge of simple cells in the cat striate cortex. *Experimental Brain Research*, 44:437–440, 1981b.
- [7] S Hecht, S Shlaer, and M H Pirenne. Energy, quanta, and vision. *J. Gen. Physiol.*, 25:819–840, 1942.
- [8] W T Newsome, K H Britten, and J A Movshon. Neuronal correlates of a perceptual decision. *Nature*, 341:52–54, 1989.
- [9] B Sakitt. Counting every quantum. *J. Physiol. (London)*, 223:131–150, 1972.
- [10] C D Salzman, K H Britten, and W T Newsome. Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346:174–177, 1990.
- [11] C D Salzman, C M Murasugi, K H Britten, and W T Newsome. Microstimulation in visual area MT: Effects on direction discrimination performance. *Journal of Neurophysiology*, 12:2331–2355, 1992.
- [12] R J Snowden, S Treue, and R A Andersen. The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Experimental Brain Research*, 88:389–400, 1992.
- [13] W R Softky and C Koch. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *Journal of Neuroscience*, 13:334–350, 1993.
- [14] D J Tolhurst, J A Movshon, and A F Dean. The statistical reliability of single neurons in cat and monkey visual cortex. *Vision Research*, 23:775–785, 1983.